AIAA-98-4733

BOEING ROTORCRAFT EXPERIENCE WITH ROTOR DESIGN AND OPTIMIZATION

Frank Tarzanin*
Darrell K. Young†

The Boeing Company
Philadelphia, Pennsylvania

Abstract

This paper reviews 12 years of progress in applying optimization to the helicopter rotor design problem. This involves multiple disciplines, multiple objective functions, a large number of design variables and irregular design space. The initial step was to develop a single interdisciplinary analysis to evaluate the objective function. By understanding the problem, approaching it incrementally and learning how to adapt optimization techniques, dramatic progress has been made. Numerous optimization techniques have been tried, including: gradient-based methods (with finite difference and automatic differentiation), biological models, surface approximations and direct search. Each of these methods had to be properly adapted to the problem. Initial progress was made using a gradient-based method along with numerous "prodding" techniques to avoid local minima. Though successful, it required extensive labor hours. In search of more efficient methods, a scaled down representative problem was defined and multiple derivative free optimization (DFO) methods were investigated. All this has led to a hybrid approach that we are currently using in rotor design.

Introduction[1]

Helicopter rotor design is a complex interdisciplinary process. Optimization was applied to this process with two major objectives in mind. First, to define a rotor with improved characteristics (lower loads, longer life, reduced weight, lower vibration, better aerodynamic performance) and second, to automate the rotor design process to reduce labor hours and design cycle time. Achieving these objectives requires many steps. As a first step, we chose to focus only on the lower vibration aspect, which is an ambitious starting point with large potential benefits. The plan was to incrementally build upon this base, adding more complexity at each step.

Historically, a major problem in the rotorcraft industry has been vibration. The primary cause of this vibration is the hub loads coming from the rotor. The transformation of the rotating vibratory hub loads into the fixed system causes a selective cancellation and addition. This results in fixed system vibratory hub loads at frequencies that are integer multiples of the number of rotor blades times the rotor speed. This frequency is represented as NP, where N is the number of blades and P represents the frequency of rotation. The fixed system vibratory vertical (Fz), longitudinal (Fx), and lateral (Fy) forces along with the roll (Mx) and pitch (My) moments at the rotor hub excite the fuselage causing vibration. The resulting vibration annoys and fatigues crew and passengers, cracks structure, and fails components and electronics. Collectively, this contributes significantly to operating cost and safety. To keep vibration reasonable, though still at undesirable levels, devices are added to most helicopters. These include absorbers in the fixed and rotating system, isolation and active force generators together with fuselage structural tuning. All these devices add cost, complexity and weight.

As a result, substantial research has been performed to reduce the inherent vibratory hub loads that cause aircraft vibration. The Ref. 1 research and unpublished wind tunnel testing showed substantial potential for reducing rotor vibratory hub loads by more than 50 percent when the blade tip was swept. However, the number of design variables, the interaction between the five different hub load components (vertical force, inplane forces and inplane moments), the real design constraints, and four years of trying convinced us that a trial-and-error, follow-the-logic approach would not work. Only computer-automated optimization could efficiently juggle all the variables and find its way through the conflicting requirements.

The objectives of this paper are to describe the steps taken thus far in the development of our rotor design tool. We will describe the various optimization techniques tried to date and show how they are being used to design low vibration rotors (LVR). We will present our experiences, including lessons learned, as applied to

---

* Manager, Dynamics and Loads
† Senior Technical Specialist

using the various optimization techniques. The results from wind tunnel model rotor tests are included to show the benefits achieved from the design process. We will also briefly describe recent activities undertaken with various researchers applying derivative free optimization (DFO) techniques to this problem. Finally, we describe some of our future plans.

The Rotor Problem

The helicopter rotor represents the classic aeroelastic problem. Figure 1 plots angle of attack versus Mach number for different blade stations. Each loop in the figure represents the travel of one blade station through one rotor rotation. The blade encounters transonic flow, stall, reverse flow (the angle of attack exceeds 180 degrees) and unsteady effects, including dynamic stall (since the blade performs multiple revolutions each second). As the blade rotates, the large changes in dynamic pressure and angle of attack result in large variations in lift. This, in turn, results in trailed and shed vortices leaving the blade as shown in Figure 2. Blades that follow run into this complex wake, referred to as non-uniform downwash, resulting in further lift variations. In addition, since the rotor blade is long and slender, there are substantial elastic deformations, including nonlinear structural dynamics such as radial shortening, Coriolis forces and bending-torsion coupling. Therefore the airloads are functions of the aircraft flight condition, the non-uniform downwash and the elastic deflections of the blade. Clearly, there is no hope of predicting rotor behavior with a loosely coupled analysis. Figure 3 shows the close coupling required to perform a complete rotor analysis.

Historically, the aerodynamics, flying qualities, dynamics and acoustics departments develop and maintain separate simulation codes for performing their tasks. The aerodynamics department is responsible for rotor performance and aero-acoustics, and developing new technology for airfoils, non-uniform downwash prediction and blade/vortex interaction. The dynamics department is responsible for rotor vibratory loads and stability, and developing aeroelastic models (blade coupled dynamic response and unsteady aerodynamics). The flying qualities department is responsible for the flight control laws and developing full aircraft trim theory. We all are trying to solve the same problem, but with different emphasis.

Each simulation has to contain most of the problem elements, but not necessarily all or the best. For aero-acoustic predictions, the blades were assumed rigid; for performance and trim predictions, approximate blade deflections were used; and for vibratory loads, simplified (quick running) downwash models were used. In the late 1980's the development of code configuration management tools (like DSEE[2] and later ClearCase[3]), increased computer power, and relentless cuts in development budgets forced a consolidation.

The aerodynamics, acoustics and dynamics departments then combined their best technology into a single interdisciplinary rotor code[4]. Code configuration management tools allowed each department to continue to develop and enhance their traditional areas of expertise and be able to utilize a simulation code that had all the best technology and was superior to any of the previous simulations. Faster computers and the proliferation of affordable workstations lessened the need to simplify portions of the theory to reduce run time and turnaround. Program options allow less rigorous, quicker running versions to be used when needed.

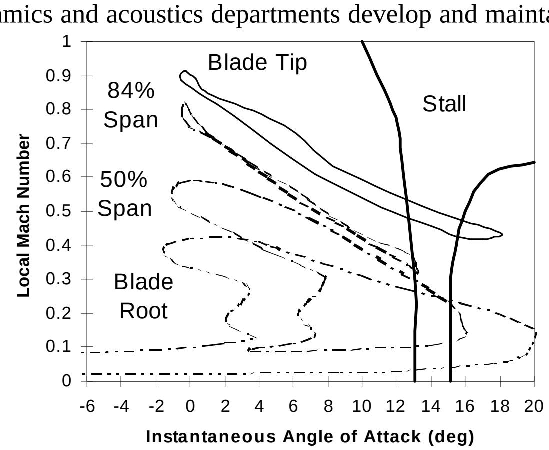

Figure 1. The challenge.

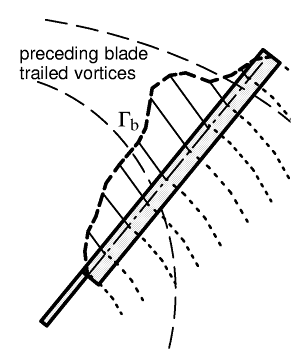

Figure 2. Trailed and shed vortices.

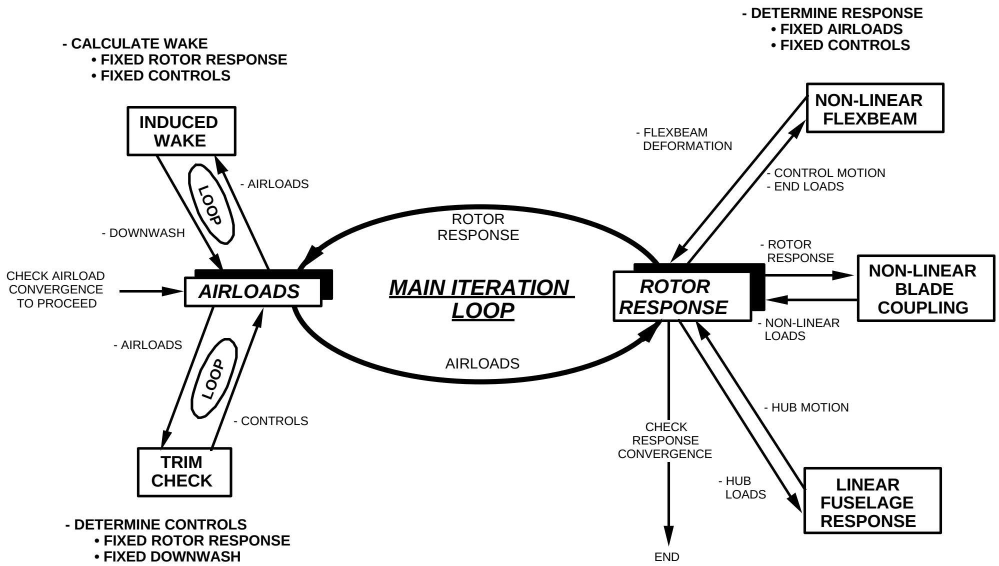

Figure 3. Rotor analysis.

Though flying qualities was not part of the initial simulation consolidation, provisions were made to make the interface with the trim model more robust. By satisfying rotor trim forces, instead of postulating control inputs, we are assured that the fuselage force and moment balance is maintained. We plan to link the combined interdisciplinary rotor analysis with the trim analysis in the future.

Therefore, our present function evaluator is a single, tightly coupled, interdisciplinary rotor analysis. An iteration method is used to satisfy compatibility among all the disciplines.

The Optimization Problem

Our first step was to prove that the optimization process worked by defining a rotor with significantly reduced vibratory hub loads, building a Mach scaled model and performing a validation test. Once validated, the second step was to obtain a better design with acceptable risk and cost. Since there are so many conflicting intangible requirements like manufacturing, total ownership cost, tolerance to variability, etc., we needed to define the design space so that the design team could find an acceptable compromise. Performing multiple point optimizations to define the local design space as a function of key variables would do this.

For our first attempt, we linked the rotor analysis, (to define the objective function), with a gradient-based optimization code, NPSOL[5]. This code is what we will refer to as the gradient optimizer throughout this paper. The rotor blade is typically modeled with 25 structural elements. There are six design variables for each element, which are listed in Table 1. Hence, there is a total of 151 available design variables (there is one extra design variable for the control system stiffness). The objective function (OF) is made up of a weighted linear sum of the five hub load components as follows:

$$OF = W_1 F_x + W_2 F_y + W_3 F_z + W_4 M_x + W_5 M_y \quad (1)$$

where

$$F_{N,D} = \sum f_{n,D,C} \quad (2)$$

$$M_{N,D} = \sum m_{n,D,C} \quad (3)$$

W_i is a coefficient for weighting the hub load components, so as to account for fuselage response due to each hub load. F and M are the fixed system hub forces and moments at N times rotor speed (where N is the number of blades), in D directions (x, y, z). Equations 2 and 3 represent the transformation of the rotating system blade root forces and moments, (f and m), in direction D, at frequency n, and flight condition C, into the fixed system forces and moments (F and M). Due to this transformation of rotating hub load components into the fixed system, there can be a shift in the frequency of one times the rotor speed. Therefore, the rotating frequency n may be at a frequency of (N-1), N, or (N+1) times the rotor speed.

This objective function formulation results in very complex design space. Since each component of the objective function has a different trend, due to a design variable change, the design space will have many peaks and valleys. Therefore, finding the lowest valley is a demanding challenge.

Table 1. Rotor design variables.

Symbol	Description
m	section mass
cg	section chordwise center of gravity
EI_β	section flap bending stiffness
EI_ζ	section lag bending stiffness
GJ	section torsion stiffness
Λ	built in sweep angle between sections
K_z	control system stiffness (only one value)

Initial Optimization Using Gradients

One major complication with the optimization process is the large demand for computer resources. Since finite differences are used to determine derivatives and there are 151 design variables, an optimization would require thousands of function evaluations. With run times of 20 to 30 minutes for each evaluation (on a HP 715/100), and questions of numerical noise, we decided to use a simple, less costly approach.

The reason the function evaluation is so computer intensive is that the airloads are a function of both the aircraft flight condition and the rotor blade elastic deflections. The elastic deflection with the largest influence on the airloads is the blade elastic twist. It was hoped that by fixing (not varying) the design variables that influence elastic twist, such as the torsional stiffness, chordwise center of gravity, aerodynamic center and chord sweep, the elastic effect on the airloads would be minimal and could be ignored. This would mean that the airloads could be assumed to be only a function of the flight condition. Therefore, the airloads could be prescribed and the rotor blade structural properties optimized to minimize the resulting vibratory hub loads. Bending-torsion coupling would still be accounted for, but changes should be minimal, so that airload changes resulting from this coupling would be small and not prevent us from finding a good dynamic response optimum.

A new, simpler function evaluator was made from the rotor blade dynamic response portion of the rotor analysis. The airloads were read into the simpler function box as a prescribed forcing function. The vibratory hub loads were calculated from the new function evaluator. Since the airloads were held fixed at the initial distribution in the optimization, the blade geometry aerodynamic configuration was also fixed. Eliminating the torsion degree of freedom from consideration reduced the number of design variables to the section mass, flap stiffness and lag stiffness at 25 blade stations. It turned out that the chord stiffness did not vary with the optimization process for reasons we do not understand. So effectively, there were only 50 design variables.

This simplification allowed the function evaluator to run in seconds. However, the optimizer still gave lackluster results. It would run through a few optimization iterations and proclaim victory, but usually the reduction in hub loads were less than twenty percent. It was clear that the gradient-based algorithm was getting stuck in local minima.

The problem is that a gradient optimizer cannot find a solution far from the initial design if the design space resembles the Rocky Mountains. There was no mechanism for a gradient-based optimizer, which follows a steepest decent, to search on the other side of a response peak, (which is perceived as first going up hill).

To resolve this problem, numerous techniques were developed to encourage the optimizer to avoid local minima. These techniques are described in more detail in the following subsections.

Different Starting Frequencies

Blade properties were changed to get different starting frequencies. By making random variations to the physical properties, new starting designs were found for the optimizer. These new designs were generated in hope of forcing the optimizer to follow a different search path. This path would either lead to the same local minimum, a different local minimum, or to the global minimum.

Large Range of Design Variable Values

By changing the range between the upper and lower bounds of the design variables, it is possible to encourage large changes in the design value. These large changes would often cause the optimizer to explore a new design space, which resulted in finding a more global minimum. Once a good solution was found, we would then squeeze the range down until we achieved a solution, which was the best compromise between hub load level and ease of manufacturing the rotor blade.

Apply Constraints Incrementally

This technique goes hand-in-hand with the large range of the design variables described above. By allowing the upper limit of a constraint, such as the total rotor weight, to be large at the start, it is possible to get into another region of the design space. Just as described above, once a good solution is found, the constraint would be squeezed to slowly force the solution into the desired design space.

Adjust Objective Function Weighting Values

Another technique that can be used to foster new solutions is varying the relationship of the weighting coefficients in equation (1). For example, if one of the hub load components is resistant to change, all of the other coefficients can be set to zero and the problem rerun. Once the optimizer has been forced into a new region, the original weighting coefficients can be used again to continue exploring the design space. Another approach is to increase or decrease the importance of a given hub load over that of the others to encourage further improvement.

Recalculate Constant Airloads

As described above, the airloads are a function of both the aircraft flight condition and the blade elastic deflections. The initial optimization process did not allow the design variables that influence elastic twist to vary and the airloads were held constant. Our practice was to verify any design solution obtained from the optimizer in our full interdisciplinary rotor analysis. While the assumption that the effect of bending-torsion coupling would not prevent us from finding a good dynamic response was true, there were times where a still better solution was found by simply updating the airloads and continuing the optimization process.

Competitions

This technique uses the different solutions, which have been generated by the above prodding techniques, to compete against each other. Individuals were given different starting designs and tried to improve the optimum. Weekly meetings would share lessons learned, eliminate the worst and continue refining the best. The comparison included items such as the objective function value, how well the design satisfied the constraints, and how realistic the blade section properties were.

Identify Related Designs

One observation is that as the number of potential designs increases, there will be promising designs with similar characteristics. By grouping similar designs together, not all of the competitions will need to be performed. This is important when time and computer resources are limited, since it is easier to eliminate a design then to perform the competitions and determine which ones to keep.

Optimizer Restart

When the optimizer terminates, the history of the objective function is reviewed. Usually this history shows an initial rapid reduction followed by a gradual leveling out. However, some times the objective function would still be declining when the optimizer would stop. There would not be the typical leveling out. When this occurred, the optimizer would be restarted and usually continued reducing the objective function. We suspect that the premature termination of the optimization is due to contamination of the Jacobian. Since the Jacobian is built from a finite difference process and uses current and historical data, noise in the numerical gradients could cause the contamination. Restarting allowed a new Jacobian to be generated and the determination of clear direction for the process to proceed.

Multiple Flight Conditions

We wanted a blade design that was robust over the whole aircraft flight regime, not just a single design condition. This is important since the rotor must operate over a wide range of airspeeds, altitudes, ambient temperatures and gross weights. By performing complete airspeed sweeps at multiple gross weights, we were able to select up to five critical flight conditions to include in the objective function. This virtually insured that the optimum would lower vibration over the whole flight regime. Typical selections included cruse at two gross weights, transition, and the corner of the flight envelope.

Initial Wind Tunnel Model

The procedure described above was developed and refined by applying it to the design of a Mach-scaled, four-bladed, fully-articulated, ten-foot diameter wind tunnel rotor which was fabricated and tested in our V/STOL wind tunnel[6]. The wind tunnel test allowed the gathering of steady-state vibratory rotor loads necessary to validate the low vibration rotor concept. The goal was to develop a rotor, which would substantially reduce the fixed system 4P vertical hub load and the fixed system 4P roll and pitch hub moments. Accomplishing this goal required the design and fabrication of two rotor blade sets — a reference rotor and a low vibration rotor. Both rotor blade sets would then be tested back-to-back in the wind tunnel.

Both rotor sets had identical blade radius, chord, twist, and airfoil shape distributions, as well as the same blade and hub attachment points. The only parameters that differed were the spanwise and chordwise distribution of the rotor mass and elastic properties. The reference rotor is a scaled version of the Boeing Model 360 experimental rotor, which flew to over 210 knots in level flight on an all-composite tandem rotor demonstrator aircraft. This rotor was designed by using the traditional approach of adjusting the rotor properties to provide adequate frequency separation from the harmonic aerodynamic forcing. The low vibration rotor was designed using the optimization techniques as described above.

A comparison of the measured normalized 4P hub loads, obtained from dynamically calibrated balances, for the reference rotor and the low vibration rotor (LVR) is shown in Figures 4 and 5 for a level flight condition corresponding to a nondimensional rotor lift, C_T'/σ of 0.07 and a nondimensional rotor propulsive force, $\overline{X}$, of 0.08. (The prime symbol is used throughout this paper to indicate a deviation from the classical definition of the marked quantities). The forces have been normalized by the nominal rotor thrust and the moments have been normalized by twice the nominal rotor thrust times the dimensional flap hinge offset. The hub loads are plotted versus rotor advanced ratio, μ', (which is defined as the free stream velocity divided by the rotor tip speed). Figure 4 shows that a 67 percent reduction was achieved in the 4P vertical hub load in the low airspeed transition region ($\mu' \approx 0.10$ or about 39 knots), and a 56 percent reduction was achieved in the high airspeed region ($\mu' \approx 0.43$ or about 183 knots) for a 3.4 percent increase in total rotor flapping weight.

Figure 5 shows that a 45 percent reduction was achieved in the measured 4P overturning hub moment in the low airspeed transition region ($\mu' \approx 0.10$), and a 77 percent reduction was achieved in the high airspeed region ($\mu' \approx 0.43$). The overturning hub moment refers to the magnitude of the vector sum of the roll and pitch hub moments.

The initial wind tunnel model design was a success. It had meet our goal of proving that the optimization process worked in defining a rotor with significantly reduced vibratory hub loads. It also showed us how labor intensive the optimization process could be. While the gradient-based approach had been successful, it left us looking for a better way of finding the global minimum. We were just getting started on the literature search when a new opportunity came along. We were asked to apply our optimization techniques to defining an advanced rotor for the CH-47 Chinook.

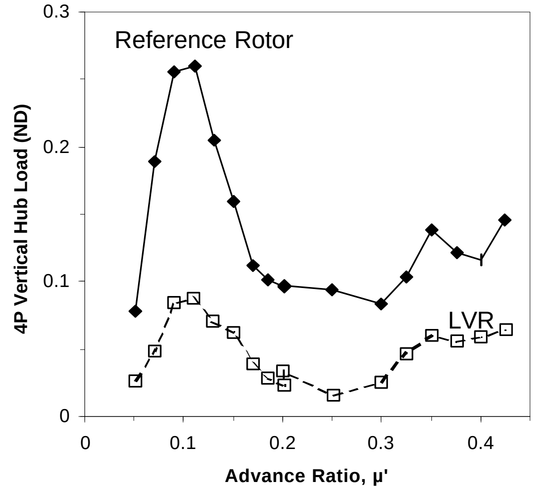

Figure 4. Measured 4P vertical hub load.

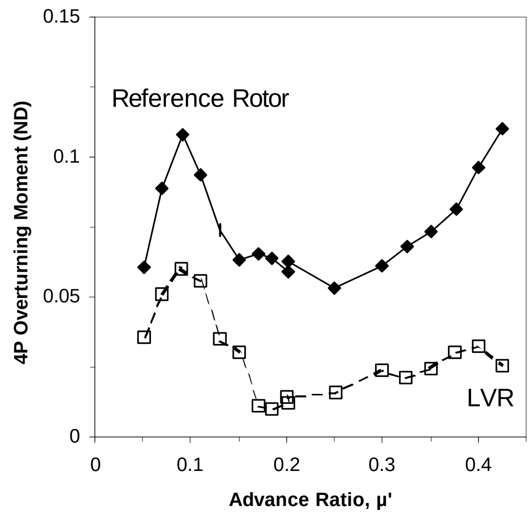

Figure 5. Measured 4P overturning hub moment.

A Real Rotor

The Mach scaled wind tunnel test results were so encouraging that funding was found to apply the optimization to an advanced CH-47 Chinook rotor. The development of a full-scale low vibration rotor was undertaken to understand and evaluate the rotor design/optimization process needed to satisfy all the full scale requirements. This included considerations like blade tie down fittings, track and balance hardware, fatigue life, tooling and manufacturing requirements.

In addition, the optimization problem was reformulated to include additional hub loads and constraints. Another wind tunnel test[7] using the previously defined LVR showed that the 8P hub loads could be measured and predicted well enough to warrant design optimization to reduce these loads as well.

The improved optimization method was applied to the design of the advanced Chinook rotor[8]. It involved working with the designers to define realistic minimum and maximum limits for each design variable. Iterating with manufacturing was required to insure that the final design was buildable. In addition, the same prodding techniques, described above, were used with the same gradient-based optimizer.

Since rotor design is a high cycle fatigue problem, the stress group periodically checked the stress/strain levels. To insure that the rotor had an infinite life, a conservative strain allowable was adapted. This strain was not to be exceeded during normal level flight. Whenever the strain was too high, the minimum blade section stiffness and weight was adjusted to lower the strain.

As the design/optimization progressed, this iteration between the optimum design, blade loads, stress and adjusted minimum stiffness and weight proved fruitless. Each time the optimizer defined a significant vibration improvement, the stress proved too high and the design constraints were adjusted. This process was increasing both our design time and cost. We either had to proceed with a less than optimum design or modify the optimization process.

Therefore, the optimization process was modified to include a maximum strain constraint. In addition, a relationship between the blade section stiffness and weight was also provided as a nonlinear constraint. This simulated the design process of strengthening the blade when the stress was too large. When additional strength was needed, the optimizer automatically added the correct weight. This made a real solution possible. As shown below, we achieved both lower vibrations, at both 4P and 8P, with a reduced blade weight.

Another issue was the determination of local design space. This would allow the design team to perform a tradeoff between total rotor weight, vibration and strain. Point optimizations were performed where the total rotor weight and strain constraints were incrementally decreased while satisfying all other constraints. By plotting the optimization results as a function of vibration versus constrained blade weight and allowable strain, the tradeoff between weight, strain, and vibration could be more clearly understood. Using this data we could choose the best compromise.

Figure 6 shows the 4P blade vibration index versus nondimensional blade weight for the design strain level and for a 30 percent larger strain level. The 4P vibration index is a normalized measure of the calculated 4P vertical force, roll moment and pitch moment, times the pilot vertical vibration response to hub loads as measured from an aircraft shake test. Hub loads from both the forward and aft rotors at 20 and 150 knots were used.

Two baseline rotors are shown. The reference model rotor (solid square) has a weighted vibration index based on calculated hub loads and is normalized to unity. The full scale Model 360 nondimensional

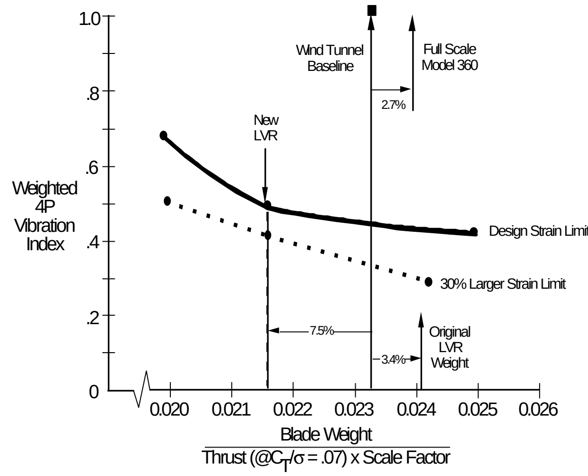

Figure 6. Weighted vibration index versus blade weight trend for two strain levels.

flapping weight is included for reference to show how close the scaled weight of the model and full scale rotors are. The nondimensional weight of the original LVR model rotor (previously discussed) is also provided for reference. This rotor is 3.4 percent heavier than the reference model rotor.

Using this design space definition, a blade weight at the knee of the curve for the lower strain was selected. This represents a flapping weight that is 7.5 percent lighter than the reference rotor and 11.9 percent lighter than the original LVR.

Observe that if desired, a new material with a higher strain allowable could be identified and qualified, or a finite blade fatigue life defined. This would result in further reduced rotor weight and/or lower vibration, but with increased cost and development risk.

The final rotor properties were Mach scaled and a wind tunnel test was performed in the same manner as the previous tests. The improved LVR used a hub with a coincident elastomeric bearing. Due to model elastomer bearing size limits, it was not possible to get the model flap hinge at the same offset as the previously described reference rotor. Therefore, to compare with the reference rotor hub loads the improved LVR measured hub moments are scaled to account for this difference.

Figures 7 and 8 show the measured normalized 4P hub loads, obtained from dynamically calibrated balances, for the reference rotor, the LVR, and the improved LVR at the same flight condition. Compared to the reference, the improved LVR shows that the 4P vertical hub load is 74 percent lower in transition and 69 percent lower in cruise for a 7.5 percent decrease in total rotor flapping weight. The 4P overturning hub moment is 88 percent lower in transition and 55 percent

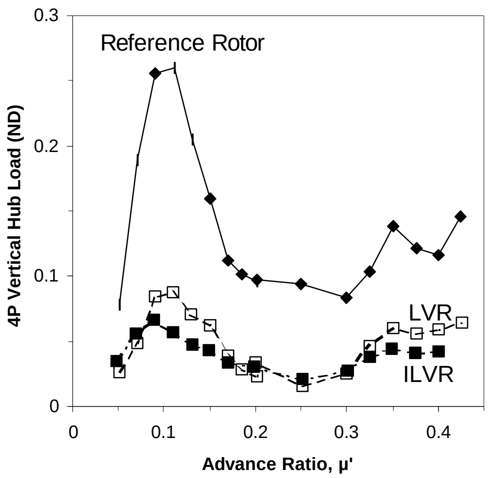

Figure 7. Measured 4P hub load for the reference rotor, LVR and ILVR.

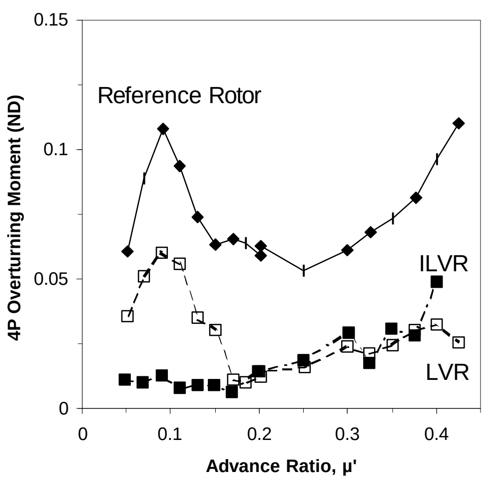

Figure 8. Measured 4P overturning hub moment for the reference rotor, LVR and ILVR.

lower in cruise. The 8P hub loads (shown in Ref. 8) had a 1 to 68 percent reduction in vertical hub load and a 53 to 79 percent reduction in overturning hub moment

Solve the Whole Problem

We achieved our objective of proving the optimization process works with a simplified function evaluator. A rotor that significantly reduced vibratory hub loads was defined and validated in a wind tunnel test.

It was now time to improve the process. This involved several enhancements. First, the full interdisciplinary rotor analysis code was used as the function box, allowing us to investigate the full aero-elastic optimization problem. Now the additional design variables that cause elastic blade twist, which leads to changes in the airloads, could be exploited by using optimization. This increased the number of design variables from 50 to 151.

The second enhancement reduced optimization turn-around time. A typical function evaluation with a prescribed set of airloads (not changing due to blade response) was a few seconds. When the full interdisciplinary rotor analysis was used (with compatibility between the blade response, airloads and rotor wake), the time increased to about 20 or 30 minutes. Hence, nearly all of our processing time would be spent evaluating the finite differences. To perform a single optimization using the full theory would require months of run time. Two developments that helped overcome this computer time problem were the continuing work-station speed increase and the use of parallel processing to evaluate each gradient independently[9].

The third enhancement focused on non-gradient based optimization methods. As pointed out above, several prodding techniques were needed to encourage a gradient-based optimization method find a global minimum. This was very time consuming and labor intensive. It was hoped that non-gradient methods would prevent getting stuck in local minima by providing a diverse set of potential optimum solutions. These potential solutions would be found by exploring the whole design space, instead of being limited to the local space of the initial design, like gradient methods. The best non-gradient designs would then be refined using gradient methods. This approach is equivalent to flying-over the Rocky Mountains to identify the most promising valleys, then sending in explorers to search for the bottom of each valley.

To help us explore the many nongradient-based optimization methods, a simple design problem was created. This problem was then given to various researchers so that they could apply their nongradient-based methods to the same rotor problem. They were also asked to use our enhanced function box evaluator.

The problem represents a three-bladed helicopter rotor with advanced planform geometry including blade tip sweep. Early optimization efforts had shown that it was more difficult to reduce vibration for a three-bladed rotor than for a four-bladed rotor. The rotor was discretized into a model consisting of 13 bays of which the 10 outboard bays had airloads applied. Normally 25 bays are used. This problem had 56 design variables which represented the level of section mass, stiffness (in flap, chord, and torsion), and chordwise center of gravity position at different blade stations along the span of the blade. Further CPU run time reductions were obtained by prescribing the rotor induced non-uniform downwash and using only two flight conditions in the function evaluation. This simplified model was chosen since it captured the main effects of the vibration problem and still had a rather short function evaluation CPU time of a few minutes per airspeed.

The objective function to be minimized was the linear combination of the weighted fixed system 3P and 6P three hub forces (Fx, Fy, Fz) and two hub moments (Mx, My). The 3P loads were weighted as being twice as important as the 6P loads and the two airspeeds were weighted equally. The only constraint limited the total rotor weight to be less than or equal to 1.685 times the nominal weight.

The methods explored were:

1) design of experiments (DOE) with response surfaces by Boeing, Seattle[10,11]
2) evolutionary programming (EP) by Boeing, Philadelphia[9]
3) parallel direct search (PDS) by Boeing Seattle, IBM, and Rice University[11,12]
4) analytical derivatives using ADIFOR by NASA, Langley[13]
5) derivative free optimization (DFO) by IBM[11,14]
6) genetic optimization (GA) with a neural net by Rensselaer Polytechnic Institute (RPI)

Table 2 compares the results obtained from each of the methods along with the results from our gradient-based method using NPSOL. Please note that the results presented here were obtained prior to the end of 1997 and that more recent results may be shown at this conference by the individual researchers. Also, we will not provide detail of how each researcher obtained their results. That too is left for the papers they will present at this conference.

Note that the Table 2 results are not global minimum and need to be refined with gradient methods (except for NPSOL and ADIFOR, which are gradient methods). Even though they are not minima, three of the derivative free methods; EP, PDS, and GA have objective function values lower than the best gradient result.

Table 2. Comparison of the resulting designs.

Design	Nondimensional Total Weight	Objective Function Value
Baseline	1.000	1.000
NPSOL	1.452	0.559
DOE	1.680	0.644
EP	1.635	0.487
PDS	1.289	0.501
ADIFOR	1.323	0.564
GA	1.685	0.512

Our experience, to date, has been that a nongradient-based method by itself is not the fastest way to reach a global minim. Because the function to be evaluated is computation intensive and many function evaluations are needed, a combination of methods is required. By automating a combined process, labor costs can be greatly reduced. For now, we have selected a hybrid approach that uses our EP method and our gradient-based method. We have chosen these two for the following reasons. First, we have both codes in house and have some, albeit limited, experience in using them. Second, the other methods are still under investigation. It is possible that multiple methods may be needed. Third, Table 2 shows that the EP method gave the best results.

The advantage of using a hybrid method is that the nongradient-based method provides a diverse set of solutions, which explore the whole design space without having to use prodding techniques. These diverse solutions are also automatically generated by the process itself and do not require labor intensive human intervention.

Recent Design Activity

Recently (last quarter of 1997), we were asked to define a replacement rotor for an existing helicopter. The new rotor would have a 12 percent larger blade chord and a 67 percent increase in blade twist but the rotor vibration could not be any higher than the existing rotor vibration. Historically, when a rotor blade has its chord increased, the section airloads increase, thereby increasing the rotor loads. In addition, increasing the blade twist also causes increased hub loads. A conventional preliminary design had been performed prior to our involvement, and the predicted vibration was substantially higher than the existing aircraft.

Using the hybrid method described above along with the complete interdisciplinary rotor analysis, the evolutionary programming method defined a promising rotor. It satisfied the vibration requirement, but was heavy and did not satisfy all the constraints. This was expected, since the initial goal was to find potential candidates, not the final design. The gradient-based method was utilized to improve this design. The weight was systematically reduced and the strain constraint applied. A dramatic improvement was made while reducing the total rotor weight by 7 percent. Figures 9 to 11 show preliminary results of the normalized hub loads for a LVR compared to the existing production rotor.

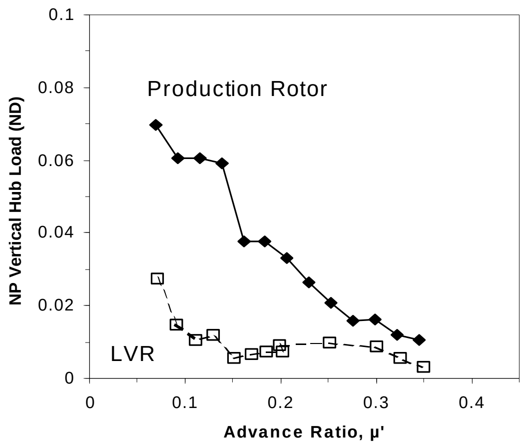

Figure 9. NP vertical hub load reduction.

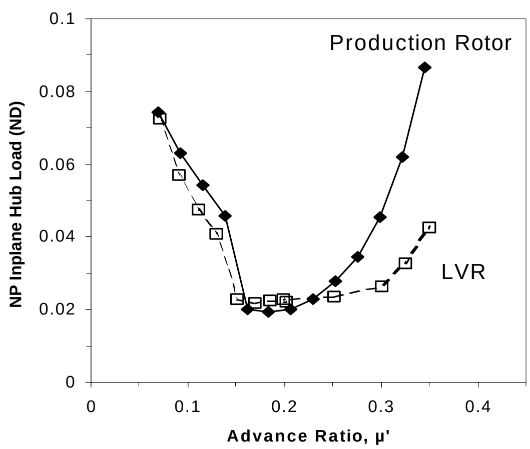

Figure 10. NP inplane hub load reduction.

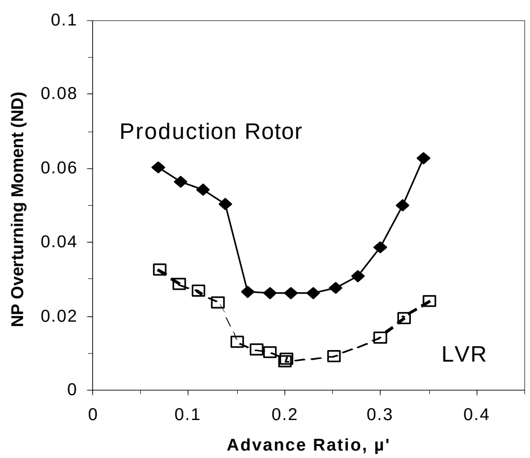

Figure 11. NP overturning hub moment reduction.

Conclusions

The rotor design problem involves a large number of design variables, interdisciplinary considerations and complex design space. The function evaluator is a single, tightly coupled, interdisciplinary, computation intensive code. Steady progress has been made towards developing an effective optimization-based rotor design process. However, the process requires excessive computer resources, long calendar time and is too labor intensive.

To date derivative free optimization (DFO) has shown the greatest promise in improving the rotor design process. By using these methods to explore the whole deign space, we are able to get many varied starting designs for refinement with our gradient-based method, and save substantial labor hours previously spent avoiding local minima. This hybrid approach has increased our confidence that a global optimum can be reached. This approach also lends itself to parallelization and we have been able to make excellent use of idle workstations.

Future Plans

Our major objectives are to define a rotor with improved characteristics (lower loads, longer life, reduced weight, lower vibration, better aerodynamic performance) and to automate the rotor design process to reduce labor hours and design cycle time. Some of the improvements described below are only notional. As we get closer to implementation, our vision will become more focused, allowing better definition of what we want to achieve.

First, we want to add rotor aerodynamic performance to the objective function. To accomplish this, more design variables and constraints must be added to the problem formulation.

Next, we want to continue investigating DFO methods. Which method is most robust (gives the best results in the least calendar time, uses less computer resources and fewest labor hours)? Are approximate methods most efficient, or are errors too large to give meaningful results? Will only using “main effects” allow substantial reductions in the number of design variables or will variable sensitivity be impossible to evaluate over the whole design space? How should the optimization control parameters be set to perform the most efficient searches? These and many other questions need to be answered.

Another improvement is the development of a method for classifying the many promising designs that result from a DFO optimization. By identifying similar designs, only the best, unique (unrelated) need to be refined with the gradient method, eliminating duplicate effort.

Data mining is another potential source of efficiency. By adding all the previously evaluated designs in a non-dimensional database, a resource can be developed for future DFO activity. Future rotor design requirements can have different emphasis on performance, loads or vibration, with different constraints. This will require a new design optimization problem. The database can be searched to rapidly define favorable designs to start the DFO process. Another application is to use the database for building response surface approximations. As more designs are investigated, the database will grow and so will the efficiencies.

The whole optimization/design process needs further automation. This may use natural language to set up the optimization, run hands off until the requested task is complete, automatically display local design space for selected parameters so intelligent tradeoffs can be made, provide status data to monitor progress, and use parametric design variable ranges and constraints for initial optimization.

We also want to improve the parallel nature of our codes. Currently we are doing most of our computation on a network of UNIX workstations. We need to improve the robustness of our controller so that if one node "crashes" (as they inevitably do), the process can continue with the remaining nodes. In addition, we want the controller to automatically search out idle computers so we can take advantage of this resource, on a noninterference basis.

Automated optimization and design are critical for the future of manufacturing in developed nations. Market forces are requiring us to design, build, and get to market faster. Reality is pushing us to reduce design cost by doing more with less. It can be done!

Acknowledgments

The authors would like to express their gratitude to Joel Hirsh for his helpful insights and knowledge of optimization techniques. His implementation of the parallel computing has been greatly valued. We would also like to thank the many researchers, too numerous to mention by name, who have worked with us in our investigation of derivative free optimization methods.

References

1. Tarzanin, F. J. and Young, D. K., "Blade Tip Sweep Effect on Hub Vibratory Loads," NASA CR 177425, Sep 1987.

2. Domain Software Engineering Environment (DSEE) Command Reference, Atria Software, Inc., Natick, Mass., 1986.

3. ClearCase User's Manuel, Atria Software, Inc., Natick, Mass., 1995.

4. Shultz, L. A., Panda, B., Tarzanin, F. J., Derham, R. C., Oh, B. K., and Dadone, L., "Interdisciplinary Analysis For Advanced Rotors — Approach, Capabilities And Status," Presented at the AHS Aeromechanics Specialists Conference, San Francisco, CA, Jan 1994.

5. Healy, M. J., "User's Guide for the SOL/NPSOL Nonlinear Programming Library — Boeing Version," Engineering Technology Applications Library Report ETA-LR-41, Boeing Computer Services, Jun 1987.

6. Young, D. K. and Tarzanin, F. J., "Structural Optimization And Mach Scale Test Validation of a Low Vibration Rotor," *Journal of American Helicopter Society*, Vol. 38,(3), Jul 1993.

7. Staley, J. A., Mathew, M. B., and Tarzanin, F. J., "Wind Tunnel Modeling of High Order Rotor Vibration," Presented at the AHS 49th Annual Forum, St. Louis, Missouri, May 1993.

8. Tarzanin, F. J., "An Improved Low Vibration Rotor," Presented at the AHS Aeromechanics Specialists Conference, Bridgeport, CT, Oct 1995.

9. Hirsh, J. E. and Young, D. K., "Evolutionary Programming Strategies with Self-Adaptation Applied to the Design of Rotorcraft using Parallel Processing," *7th Annual Conference on Evolutionary Programming*, San Diego, CA, March 25-27, 1998, Springer-Verlag.

10. Booker, A. J., "DACE - Design and Analysis of Computer Experiments," Presented at the 7th AIAA/USAF/NASA /ISSOMO Multidisciplinary Analysis and Optimization Symposium, St. Louis, Missouri, Sep 2-3, 1998.

11. Booker, A. J., Dennis Jr., J. E., Frank, P. D., Serafini, D. B., Torczon, V., and Trosset, M. W., "A Rigorous Framework for Optimization of Expensive Functions by Surrogates," Technical Report SSGTECH-98-095, Boeing Shared Services Group, Applied Research and Technology, March 1998.

12. Dennis Jr., J. E. and Serafini, D. B., "Model Management," Presented at the 7th AIAA/USAF/NASA/ISSOMO Multidisciplinary Analysis and Optimization Symposium, St. Louis, Missouri, Sep 2-3, 1998.

13. Walsh, J. L., Young, K. C., Tarzanin, F. J., Hirsh, J. E., and Young, D. K., "Optimization Issues with Complex Rotorcraft Comprehensive Analysis," Presented at the 7th AIAA/USAF/NASA/ISSOMO Multidisciplinary Analysis and Optimization Symposium, St. Louis, Missouri, Sep 2-3, 1998.

14. Conn, A. and Mints, K., "Derivative Free Optimization," Presented at the 7th AIAA/USAF/NASA/ISSOMO Multidisciplinary Analysis and Optimization Symposium, St. Louis, Missouri, Sep 2-3, 1998.